

IN THE CLAIMS:

1 1. (Currently Amended) A method for a storage operating system implemented in a
2 storage system to optimize an amount of readahead data retrieved from
3 a data container of the storage system, the method comprising:
4 maintaining a plurality of files;
5 maintaining, for a selected file of the plurality of files, a plurality of readset data
6 structures, each readset data structure holding a plurality of factors for a selected
7 readstream;
8 receiving a client read request for a particular read stream at the storage system;
9 locating a readset data structure for the particular read stream;
10 determining whether the storage operating system is permitted to retrieve
11 readahead data from the data container in response to the received client read request;
12 if it is determined that the storage operating system is permitted to retrieve
13 readahead data from the data container, performing the steps of [:] ,
14 (i) selecting an amount of readahead data to retrieve from the data container,
15 based on [a] the plurality of factors stored within the readset data structure; and
16 (ii) retrieving the selected amount of readahead data from the data container;
17 and
18 (iii) determining if the readahead data structure meets a criteria for being
19 updated, and if the readahead data structure meets the criteria then updating the
20 readahead data structure .

1 2. (Original) The method of claim 1, wherein the data container is a file, directory,
2 vdisk or lun.

1 3. (Original) The method of claim 1, wherein the storage operating system is
2 determined to be permitted to retrieve readahead data from the data container when the
3 client-requested data extends the read stream past a predetermined next readahead value.

1 4. (Original) The method of claim 3, wherein the predetermined next readahead
2 value is stored in a readset data structure associated with the read stream.

1 5. (Original) The method of claim 3, wherein the predetermined next readahead
2 value is updated based on a percentage of the selected amount of readahead data.

1 6. (Previously Presented) The method of claim 1, wherein a read-access style
2 associated with the data container is one of the plurality of factors used to select the
3 amount of readahead data.

1 7. (Original) The method of claim 6, wherein the selected amount of readahead data
2 equals zero if the read-access style corresponds to a random read-access style.

1 8. (Previously Presented) The method of claim 1, wherein a number of client read
2 requests processed in the read stream is one of the plurality of factors used to select the
3 amount of readahead data.

1 9. (Original) The method of claim 8, wherein the number of client read requests
2 processed in the read stream is stored as a count value in a readset data structure
3 associated with the read stream.

1 10. (Previously Presented) The method of claim 1, wherein the amount of client-
2 requested data is one of the plurality of factors used to select the amount of readahead
3 data.

1 11. (Original) The method of claim 10, wherein the selected amount of readahead
2 data is set equal to a predetermined upper limit for large amounts of client-requested data.

1 12. (Original) The method of claim 1, wherein the selected amount of readahead data
2 is doubled if the number of client read requests processed in the read stream is greater
3 than a first threshold value.

1 13. (Original) The method of claim 1, wherein the client-requested data is identified
2 as read-once data when either (i) the number of client read requests processed in the read
3 stream is greater than a second threshold value or (ii) a set of metadata associated with
4 the read stream indicates that the client-requested data is read-once data.

1 14. (Original) The method of claim 1, wherein the selected amount of readahead data
2 is stored in one or more buffers enqueued on a flush queue, the flush queue being
3 configured to reuse buffers after a predetermined period of time.

1 15. (Original) The method of claim 14, wherein the predetermined period of time
2 equals two seconds.

1 16. (Currently Amended) An apparatus configured to implement a storage operating
2 system that optimizes an amount of readahead data retrieved from a data container of the
3 apparatus, the apparatus comprising:

4 means for maintaining a plurality of files;

5 means for maintaining, for a selected file of the plurality of files, a plurality of
6 readset data structures, each readset data structure holding a plurality of factors for a
7 selected readstream;

8 means for receiving a client read request for a particular read stream;

9 means for locating a readset data structure for the particular read stream;

10 means for determining whether the storage operating system is permitted to
11 retrieve readahead data from the data container in response to the received client read
12 request;

13 means for selecting an amount of readahead data to retrieve from the data
14 container based on the [a] plurality of factors stored within the readset data structure;
15 and

16 means for retrieving the selected amount of readahead data from the data
17 container;

18 means for determining if the readahead data structure meets a criteria for being
19 updated, and if the readahead data structure meets the criteria then updating the
20 readahead data structure .

- 1 17. (Original) The apparatus of claim 16, wherein the data container is a file,
- 2 directory, vdisk or lun.

- 1 18. (Original) The apparatus of claim 16, wherein the storage operating system is
- 2 determined to be permitted to retrieve readahead data from the data container when the
- 3 client-requested data extends the read stream past a predetermined next readahead value.

- 1 19. (Original) The apparatus of claim 18, further comprising means for updating the
- 2 predetermined next readahead value based on a percentage of the selected amount of
- 3 readahead data.

- 1 20. (Previously Presented) The apparatus of claim 16, wherein the plurality of factors
- 2 used to select the amount of readahead data includes at least one of:
 - 3 (i) the amount of client-requested data,
 - 4 (ii) a number of client read requests processed in the read stream, and
 - 5 (iii) a read-access style associated with the data container.

- 1 21. (Original) The apparatus of claim 16, wherein the selected amount of readahead
- 2 data is doubled if the number of client read requests processed in the read stream is
- 3 greater than a first threshold value.

1 22. (Currently Amended) A storage system configured to optimize an amount of
2 readahead data retrieved from a data container of the storage system, the storage system
3 comprising:
4 a network adapter for receiving a client read request for a particular read stream;
5 a memory configured to store instructions for implementing a storage operating
6 system that performs the steps of:
7 locating a readset data structure for the particular read stream; and
8 selecting, in response to the readset data structure for the particular read stream, a
9 selected amount of readahead data from the data container;
10 retrieving the selected amount of readahead data from the data container.

1 23. (Original) The storage system of claim 22, wherein the data container is a file,
2 directory, vdisk or lun.

1 24. (Original) The storage system of claim 22, wherein the storage operating system
2 is determined to be permitted to retrieve readahead data from the data container when the
3 client-requested data extends the read stream past a predetermined next readahead value.

1 25. (Original) The storage system of claim 24, wherein the predetermined next
2 readahead value is updated based on a percentage of the selected amount of readahead
3 data.

1 26. (Previously Presented) The storage system of claim 22, wherein the plurality of
2 factors used to select the amount of readahead data includes at least one of:
3 (i) the amount of client-requested data,
4 (ii) a number of client read requests processed in the read stream, and
5 (iii) a read-access style associated with the data container.

1 27. (Original) The storage system of claim 22, wherein the selected amount of
2 readahead data is doubled if the number of client read requests processed in the read
3 stream is greater than a first threshold value.

1 28. (Currently Amended) A computer-readable media comprising instructions for
2 execution in a processor for the practice of a method for a storage operating system
3 implemented in a storage system to optimize an amount of readahead data retrieved from
4 a data container of the storage system, the method comprising:
5 maintaining a plurality of files;
6 maintaining, for a selected file of the plurality of files, a plurality of readset data
7 structures, each readset data structure holding a plurality of factors for a selected
8 readstream;
9 receiving a client read request for a particular read stream at the storage system;
10 locating a readset data structure for the particular read stream;
11 determining whether the storage operating system is permitted to retrieve
12 readahead data from the data container in response to the received client read request;

13 if it is determined that the storage operating system is permitted to retrieve
14 readahead data from the data container, performing the steps of:
15 (i) selecting an amount of readahead data to retrieve from the data container
16 based on [a] the plurality of factors stored within the readset data structure; and
17 (ii) retrieving the selected amount of readahead data from the data container;
18 (iii) determining if the readahead data structure meets a criteria for being
19 updated, and if the readahead data structure meets the criteria then updating the
20 readahead data structure .

1 29. (Original) The computer-readable media of claim 28, wherein the data container
2 is a file, directory, vdisk or lun.

1 30. (Previously Presented) The method of claim 1, wherein the retrieved readahead
2 data is stored in one or more buffers, the buffers containing a flush queue, the flush queue
3 being configured to reuse buffers after a predetermined period of time.

1 31. (Previously Presented) The method of claim 30, wherein the read stream
2 corresponds to a read-once data transfer and data retrieved from the data container is
3 stored in the flush queue.

1 32. (Previously Presented) The method of claim 30, wherein the retrieved readahead
2 data is stored in the flush queue.

1 33. (Previously Presented) The method of claim 30, wherein one or more buffers
2 accessed from the flush queue are re-enqueued on a normal queue.

1 34. (Currently Amended) A method for optimizing readahead data retrieved from
2 a data container of a storage system, the method comprising:
3 maintaining a plurality of files;
4 maintaining, for a selected file of the plurality of files, a plurality of readset data
5 structures, each readset data structure holding a plurality of factors for a selected
6 readstream;
7 receiving a client read request for a particular read stream;
8 locating a readset data structure for the particular read stream;
9 selecting an amount of readahead data based on the indicated amount of client-
10 requested data and the a corresponding readset data structure; and
11 retrieving the selected amount of readahead data from the data container.

1 35. (Previously Presented) The method of claim 34, wherein the selected amount of
2 readahead data is set equal to a multiple of a predetermined amount, and wherein the
3 multiple is associated with the amount of client-requested data.

1 36. (Previously Presented) The method of claim 34, wherein the selected amount of
2 readahead data is set equal to a multiple of the amount of client-requested data.

1 37. (Previously Presented) The method of claim 36, further comprising the step of
2 rounding the selected amount of readahead data to the size of a data block.

1 38. (Previously Presented) The method of claim 34, wherein the selected amount of
2 readahead data is set equal to a predetermined upper limit.

1 39. (Currently Amended) A method for optimizing readahead data retrieved from a
2 data container of a storage system, the method comprising:

3 maintaining a plurality of files;
4 maintaining, for a selected file of the plurality of files, a plurality of readset data
5 structures, each readset data structure holding a plurality of factors for a selected
6 readstream;

7

8 receiving a client read request for a particular read stream at the storage system;
9 locating a readset data structure for the particular read stream;
10 selecting for the read stream an amount of readahead data based on a read-access
11 style associated with the data container and ~~the a~~ corresponding readset data structure;
12 and
13 retrieving the selected amount of readahead data from the data container.

1 40. (Previously Presented) The method of claim 39, wherein the selected amount of
2 readahead data equals zero if the read-access style corresponds to a random read-access
3 style.

1 41. (Currently Amended) A method for optimizing readahead data retrieval from
2 a data container of a storage system associated with a number of storage devices, the
3 method comprising:

4 maintaining a plurality of files;
5 maintaining, for a selected file of the plurality of files, a plurality of readset data
6 structures, each readset data structure holding a plurality of factors for a selected
7 readstream;
8 receiving a client read request for a particular read stream at the storage system;
9 locating a readset data structure for the particular read stream;
10 selecting an amount of readahead data based on ~~the~~ a number of storage devices
11 and ~~the~~ a corresponding readset data structure; and
12 retrieving the selected amount of readahead data from the data container.

1 42. (Previously Presented) The method of claim 41, wherein the step of selecting an
2 amount of readahead data further comprises:

3 determining whether a flag is associated with the read stream, the flag indicating
4 that the storage system is associated with more than a predetermined number of storage
5 devices; and
6 in response to determining whether the flag is associated, selecting the amount of
7 readahead data.

1 43. (Previously Presented) The method of claim 41, wherein the storage devices
2 comprise one or more disks.

1 44. (Currently Amended) A method for optimizing readahead data retrieval in a
2 storage system, the method comprising:
3 maintaining a plurality of files;
4 maintaining, for a selected file of the plurality of files, a plurality of readset data
5 structures, each readset data structure holding a plurality of factors for a selected
6 readstream;
7 receiving a client read request for a particular read stream at the storage system;
8 locating a readset data structure for the particular read stream;
9 selecting an amount of readahead data based on a plurality of factors stored
10 within ~~the a~~ corresponding readset data structure; and
11 retrieving the selected amount of readahead data from a data container.

1 45. (Previously Presented) The method of claim 44, wherein the retrieved readahead
2 data is stored in one or more buffers, the buffers containing a flush queue, the flush queue
3 being configured to reuse buffers after a predetermined period of time.

1 46. (Previously Presented) The method of claim 45, wherein the read stream
2 corresponds to a read-once data transfer and data retrieved from the data container is
3 stored in the flush queue.

1 47. (Previously Presented) The method of claim 45, wherein the retrieved readahead
2 data is stored in the flush queue.

1 48. (Previously Presented) The method of claim 45, wherein one or more buffers
2 accessed from the flush queue are re-enqueued on a normal queue.

1 49. (Currently Amended) A system for optimizing readahead data retrieval in a
2 storage system, the system comprising:
3 maintaining a plurality of files;
4 maintaining, for a selected file of the plurality of files, a plurality of readset data
5 structures, each readset data structure holding a plurality of factors for a selected
6 readstream;
7 means for receiving a client read request for a particular read stream at the storage
8 system;
9 means for locating a readset data structure for the particular read stream;
10 means for selecting an amount of readahead data based on a plurality of factors
11 stored within ~~the~~ a corresponding readset data structure; and
12 means for retrieving the selected amount of readahead data from a data container.

1 50. (Previously Presented) The system of claim 49, wherein the retrieved readahead
2 data is stored in one or more buffers, the buffers containing a flush queue, the flush queue
3 being configured to reuse buffers after a predetermined period of time.

1 51. (Previously Presented) The system of claim 50, wherein the read stream
2 corresponds to a read-once data transfer and data retrieved from the data container is
3 stored in the flush queue.

1 52. (Previously Presented) The system of claim 50, wherein the retrieved readahead
2 data is stored in the flush queue.

1 53. (Previously Presented) The system of claim 50, wherein one or more buffers
2 accessed from the flush queue are re-enqueued on a normal queue.

1 54. (Currently Amended) A method, comprising:
2 maintaining a plurality of files;
3 maintaining, for a selected file of the plurality of files, a plurality of readset data
4 structures, each readset data structure holding a plurality of factors for a selected
5 readstream;
6 receiving a plurality of client read requests for a particular read stream at a storage
7 system;
8 locating a readset data structure for the particular read stream;
9 selecting an amount of readahead data to retrieve from one or more data
10 containers based on a plurality of factors stored within the readset data structure;
11 retrieving the selected amount of readahead data from the data container;
12 processing one or more of the plurality of client read requests; and

13 adjusting, as client requests are processed, the plurality of factors stored within
14 the readset data structure associated with each read stream to optimize the amount of
15 readahead data that is cached for each read stream.

1 55. (Previously Presented) The method of claim 54, further comprising:
2 determining whether the storage operating system is permitted to retrieve
3 readahead data from the one or more data containers in response to each received client
4 read request.

1 56. (Previously Presented) The method of claim 54, wherein the one or more data
2 containers are at least one of a file, a directory, a vdisk or a lun.

1 57. (Previously Presented) The method of claim 55, wherein the storage operating
2 system is determined to be permitted to retrieve readahead data from the one or more data
3 containers when the client-requested data extends the read stream past a predetermined
4 next readahead value.

1 58. (Previously Presented) The method of claim 57, wherein the predetermined next
2 readahead value is stored in a readset data structure associated with the read stream.

1 59. (Previously Presented) The method of claim 57, wherein the predetermined next
2 readahead value is updated based on a percentage of the selected amount of readahead
3 data.

1 60. (Previously Presented) The method of claim 54, wherein a read-access style
2 associated with the one or more data containers is one of the plurality of factors used to
3 select the amount of readahead data.

1 61. (Previously Presented) The method of claim 60, wherein the selected amount of
2 readahead data equals zero if the read-access style corresponds to a random read-access
3 style.

1 62. (Previously Presented) The method of claim 54, wherein a number of client read
2 requests processed in the read stream is one of the plurality of factors used to select the
3 amount of readahead data.

1 63. (Previously Presented) The method of claim 62, wherein the number of client read
2 requests processed in the read stream is stored as a count value in a readset data structure
3 associated with the read stream.

1 64. (Previously Presented) The method of claim 54, wherein the amount of client-
2 requested data is one of the plurality of factors used to select the amount of readahead
3 data.

1 65. (Previously Presented) The method of claim 64, wherein the selected amount of
2 readahead data is set equal to a predetermined upper limit for large amounts of client-
3 requested data.

1 66. (Previously Presented) The method of claim 54, wherein the selected amount of
2 readahead data is doubled if the number of client read requests processed in the read
3 stream is greater than a first threshold value.

1 67. (Previously Presented) The method of claim 55, wherein the client-requested data
2 is identified as read-once data when either (i) the number of client read requests
3 processed in the read stream is greater than a second threshold value or (ii) a set of
4 metadata associated with the read stream indicates that the client-requested data is read-
5 once data.

1 68. (Previously Presented) The method of claim 54, wherein the selected amount of
2 readahead data is stored in one or more buffers enqueued on a flush queue, the flush
3 queue being configured to reuse buffers after a predetermined period of time.

1 69. (Currently Amended) A method for optimizing readahead data retrieval for a
2 storage system, the method comprising:
3 maintaining a plurality of files;
4 maintaining, for a selected file of the plurality of files, a plurality of readset data
5 structures, each readset data structure holding a plurality of factors for a selected
6 readstream;
7 receiving a client read request for a particular read stream at the storage system;
8 locating a readset data structure for the particular read stream;

9 selecting an amount of readahead data in response to the a corresponding readset
10 data structure; and
11 retrieving the selected amount of readahead data from a data container.

1 70. (Previously Presented) The method of claim 69, further comprising:
2 allocating more readsets for the file in response to processing one or more client
3 “write” requests to the file.

1 71. (Currently Amended) A storage system, comprising:
2 an operating system to maintain a plurality of files, and to maintain, for a selected
3 file of the plurality of files, a plurality of readset data structures, each readset data
4 structure holding a plurality of factors for a selected readstream;
5 a network adapter to receive a client read request at the storage system for client-
6 requested data stored in a file;
7 a plurality of readset data structures associated with the client-requested data;
8 an operating system to establish a read stream corresponding to each readset data
9 structure of the plurality of readset data structures;
10 a process to select an amount of readahead data in response to a corresponding
11 readset data structure of the plurality of readset data structures; and
12 a storage adapter to retrieve the selected amount of readahead data from the data
13 container; and

14 the operating system to determine if the readset data structure meets a criteria for
15 being updated, and if the readahead data structure meets the criteria, then updating the
16 readahead data structure.

1 72. (Previously Presented) The storage system of claim 71, further comprising:
2 the operating system to allocate more readsets for the file in response to
3 processing one or more client “write” requests to the file.